IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF MASSACHUSETTS

SKYLINE SOFTWARE SYSTEMS, INC.,

Plaintiff,

v.

KEYHOLE, INC., and
GOOGLE INC.

Defendants.

CIVIL ACTION NO. 06-10980 DPW

DECLARATION OF PROFESSOR STEVEN K. FEINER, Ph.D., IN SUPPORT OF DEFENDANTS' OPPOSITION TO PLAINTIFF SKYLINE SYSTEMS, INC.'S MOTIONS FOR SUMMARY JUDGMENT OF INFRINGEMENT AND VALIDITY

[PUBLIC REDACTED VERSION]

Nelson G. Apjohn (BBO No. 020373)
NUTTER McCLENNEN & FISH LLP
World Trade Center West
155 Seaport Boulevard
Boston, MA 02210
(617) 439-2000
Fax: (617) 310-9000

Attorneys for Defendants and Counterclaimants
KEYHOLE, INC. and GOOGLE INC.

Of Counsel
Darryl M. Woo, admitted *pro hac vice*
Michael J. Sacksteder, admitted *pro hac vice*
Heather N. Mewes, admitted *pro hac vice*
Carolyn Chang, admitted *pro hac vice*
Saundra L.M. Riley, admitted *pro hac vice*
FENWICK & WEST LLP
Embarcadero Center West
275 Battery Street
San Francisco, CA 94111
Tel. (415) 875-2300
Fax (415) 281-1350
email: dwoo@fenwick.com

I, Steven K. Feiner, declare as follows:

1. I make this declaration in support of Google's opposition to Skyline's Motions for Summary Judgment of Infringement and Validity. In this declaration, I provide my expert opinions concerning certain issues in this lawsuit relating to United States Patent No. 6,496,189 ("the '189 patent").

2. On January 19, 2007, I submitted a declaration in support of Google's Motions for Summary Judgment of Noninfringement and Validity ("Feiner SJ Decl."), wherein I addressed many of the issues raised by Skyline in its motions. Accordingly, I incorporate this declaration by reference.

3. My qualifications are stated more fully in my *curriculum vitae,* a true and correct copy of which was attached as Exhibit A to the Feiner SJ Decl.

4. I received a Ph.D. in Computer Science from Brown University in 1987. I received an A.B. degree in music from Brown University in 1973.

5. I am presently a Professor of Computer Science at Columbia University, a position that I have held for twenty years. I have been a Full Professor since January 2000. Prior to that, I was an Associate Professor of Computer Science at Columbia University from January 1991 until December 1999, and an Assistant Professor from September 1985 to December 1990. Prior to joining the faculty of Columbia University in September 1985, I was a Research and Teaching Assistant in the Department of Computer Science at Brown University from September 1977 until August 1985.

6. At Columbia University, I direct the Columbia University Computer Graphics and User Interfaces Laboratory, and teach both graduate and undergraduate students in computer graphics and user interfaces courses. I advise Computer Science doctoral candidates, primarily

in the field of computer graphics and user interfaces. I am an active academic researcher, whose areas of research include knowledge-based design of graphics and multimedia, user interfaces, virtual reality and augmented reality, wearable computing, animation, hypermedia, and visualization.

7. I am coauthor of *Computer Graphics: Principles and Practice, Second Edition,* Addison-Wesley, 1990 ("*Computer Graphics*"), an authoritative and frequently cited academic computer graphics text. I am also a coauthor of *Introduction to Computer Graphics*, Addison-Wesley, 1993, and *Computer Graphics: Principle and Practice, Second Edition in C*, Addison-Wesley, 1996. As indicated on my *curriculum vitae*, I am the author and coauthor of over thirty journal papers, over seventy conference papers, and numerous other workshop papers, books and book chapters, editorials and other publications on computer graphics and user interfaces. I have been an Associate Editor of *ACM Transactions on Graphics* and *ACM Transactions on Information Systems*, and have been on the editorial boards of *IEEE Transactions on Visualization and Computer Graphics*, and *Virtual Reality*. I am a frequent invited speaker on computer graphics and user interfaces at institutions such as Princeton University, the Massachusetts Institute of Technology, and Carnegie Mellon University. In addition, I have given invited talks at numerous conferences and workshops, including ones related to Geographic Information Systems ("GIS"), such as *GIScience 2002*, the Advanced Research and Development Activity *Geospatial Intelligence Information Visualization Researchers Meeting 2003* and *GIS Planet 2005*. In 1991, I received an Office of Naval Research Young Investigator Award.

8. I am a named inventor on an issued United States patent relating to computer graphics, entitled "Worlds-within-worlds nested display and interaction system and method"

(U.S. Pat. No. 5,524,187).

9. I have reviewed the ’189 patent and its relevant prosecution history and am familiar with this patent, its claims, and the background technology.

10. The ’189 patent uses concepts, nomenclature, designs, and systems from the computer graphics art that should be understood in this context. In my opinion, one of ordinary skill in the art relevant to the subject matter of the ’189 patent at the time the application for the patent was filed would be a person with a bachelor’s degree in Computer Science, including at least one course in computer graphics, or with academic or work experience equivalent to that level of education.

I. <u>SKYLINE’S MOTION FOR SUMMARY JUDGMENT OF INFRINGEMENT</u>

11. I understand that Skyline asserts in its motion for summary judgment of infringement that claims 1 and 12 of the ’189 patent are infringed by the accused Google Earth products. My expert reports of August 10, 2006 and December 22, 2006 (Feiner SJ Decl., Exs. B and C, respectively) provide my opinions explaining fully why I disagree with this assertion. In addition, I submitted a declaration in support of Google’s motions for summary judgment of noninfringement and invalidity (“Feiner SJ Decl.”) which includes my opinions on why the Google Earth products do not infringe the asserted claims of the ’189 patent, and which I incorporate herein by reference. In this declaration, I address some of the specific arguments raised by Skyline in its motion for summary judgment of infringement, and provide my expert opinion as to why certain limitations of claims 1 and 12 of the ’189 patent are not met by the accused Google Earth products.

A. *All Asserted Claims: “receiving from the renderer one or more coordinates in the terrain along with indication of a respective resolution level”*

12. The Court construed a “renderer” as a “software and/or hardware object that

performs at least the following functions: (1) determining and providing to another object the required coordinates in the terrain along with a respective resolution level; (2) receiving the data blocks corresponding to the specified coordinates; and (3) using the received data blocks to display a three-dimensional image." Chang Decl., Ex. 4 at 26-32. The Court construed "coordinates in the terrain" as a "set of numerical values that identifies a particular location in the terrain." *Id.* at 19-23.

13. There is no single object in the Google Earth client software that performs the first function of the claimed "renderer"; nothing determines and provides the required coordinates in the terrain along with a respective resolution level to another object in the system.

14. [redacted] does not constitute "coordinates in the terrain," as it does not identify "a particular location *in the terrain*."

[redacted]

[redacted]

[redacted]

17. Skyline also fails to establish that coordinates along with *a respective resolution level* are provided to another object. [redacted] is not "the identification of the particular resolution level of the data required."

[redacted]

19. Skyline also fails to establish that the purported coordinates in the terrain are provided to another object *along with* [redacted]

[redacted].

20. Skyline also argues that the three required functions of the claimed "renderer" do not need to be performed by a single object or logical entity. *Id.* at 25. The Court's construed "renderer" to be a "software and/or hardware object." Chang Decl., Ex. 4 at 32. It then clarified that "object" was used "in the more general sense of a 'thing,' which could be a software object, but need not be." *Id.*, Ex. 5 at 9. I understand this to mean that the "renderer" of the '189 patent must be a single thing.

21. The Court also construed "receiving from the renderer" to mean "something distinct from the renderer receiving from the renderer." *Id.* at 10.

22. Skyline's motion claims that I agree that all of the source code in a routine found in a file would not necessarily be part of a module. Skyline's Infringement Motion at 26. Skyline misstates my deposition testimony. [redacted].

B. *All Asserted Claims: "downloading ... if the provided block from the local memory is not at the indicated resolution level"*

23. The Court construed "downloading ... if the provided block from the local

memory is not at the indicated resolution level" to mean "downloading … upon some determination that the block provided from local memory is not at the indicated resolution level." Chang Decl., Ex. 5 at 12.

24. Skyline's motion argues that visual observation of the Google Earth software program demonstrates that this limitation is met. Skyline's Infringement Motion at 21-22. However, one cannot determine from mere observation whether Google Earth makes a determination that the block provided from local memory is not at the indicated resolution level.

25. There are many ways a software program could display images at a blurry lower resolution and then sharpen to a higher resolution image without making the determination required by the '189 patent.

C. ***All Asserted Claims: "data blocks belonging to a hierarchical structure"***

26. The Court construed "data blocks belonging to a hierarchical structure" to mean "data blocks that are organized into multiple levels of resolution, whereby each level contains data blocks at the same resolution, and each successive level contains data blocks of a higher resolution than those in the preceding level." Chang Decl., Ex. 4 at 15.

[redacted]

[redacted]

[redacted]

29. In its motion for summary judgment of infringement, Skyline cites the '189 patent at col. 8:38-58 to argue that vector data has no impact on the hierarchical structure. Skyline's Infringement Motion at 16. Column 8, lines 38-58 of the '189 patent, however, does not state that vector data is not part of the claimed hierarchical structure. In fact, the cited portion of the '189 patent specification explains that the vector data is overlaid and included in each data block, and would therefore be a part of the hierarchical structure. In that sense, the vector data is associated with data of a particular resolution. The fact that vector data may be optional does not affect the fact that it must be organized into the claimed "hierarchical structure" if it is present.

30. As Skyline's expert acknowledges, [redacted] As a result, Google Earth does not have "data blocks belonging to a hierarchical structure."

II. SKYLINE'S MOTION FOR SUMMARY JUDGMENT OF VALIDITY

31. I understand that Skyline asserts in its motion that no prior art reference anticipates or renders obvious claims 1 and 12 of the '189 patent. My expert report (Feiner SJ

Decl., Ex. D) provides my opinions explaining fully why I disagree with this assertion. In this declaration, however, I address some of the specific arguments raised by Skyline, and also provide my expert opinion as to why at least the following prior art references anticipate or render obvious claims 1 and 12 of the ’189 patent: (1) the public use of the TerraVision application; (2) the MAGIC Final Report; (3) the MAGIC IEEE Article; (4) the Mayer patent; (5) the T_Vision Project materials on the SIGGRAPH ’95 Multimedia CD-ROM; (6) the public use of the T_Vision application; and (7) the Migdal patent in combination with the Cosman article.

32. In its motion, Skyline makes a number of general claims about the ’189 patent and the prior art. However, Skyline’s motion fails to correctly or consistently apply the Court’s claim constructions in this case. For example, in its motion, Skyline equates “terrain” and “elevation data.” *See* Skyline Validity Motion at 1 n.2 (“For the purposes of discussing the prior art, ‘terrain’ and elevation data are used interchangeably.”). This is not correct. The Court has construed “terrain” in the ’189 patent to mean the “surface features of an area of land, an object, or a material, including color, elevation, and existing objects or structures on the land, object or material.” Chang Decl., Ex. 4 at 19. Thus, while “terrain” includes elevation data, it also includes “color” and “existing objects or structures on the land, object or material.” Both “color” and “existing objects or structures on the land, object or material” describe image data in the Court’s construction. In fact, in the hierarchical data structure of the preferred embodiment described in the ‘583 patent, tiles are represented as array of terrain units, where each “terrain unit T, typically comprises at least color and elevation attributes T_e.” ‘583 patent, col. 6:22-23. Skyline also applies different interpretations to find a “renderer” and a “hierarchical structure” in Google Earth, but no “renderer” or “hierarchical structure” in some of the prior art references.

This is also not correct. If the same interpretation is used, then Skyline must concede either that Google Earth does not infringe or that the prior art does in fact disclose these limitations.

33. Moreover, image or "color" data plays a key role in rendering *three-dimensional* terrain. Both image data *and* elevation data are used to render three-dimensional terrain, both in the '189 patent and in the prior art. *See, e.g.,* '189 patent, col. 8:32-37; Chang Decl., Ex. 21 at GOOG 359 & Ex. 22 at GOOG 349-50; Mewes Decl., Ex. 2. For example, in the TerraVision application, image data was mapped to a digital elevation model to render three-dimensional terrain. Chang Decl., Ex. 22 at GOOG 350. Image data was also used to create perspective as illustrated in Figure 3 of the MAGIC IEEE Article:

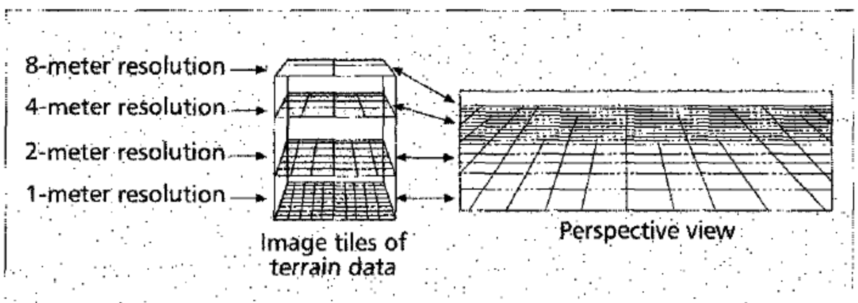

■ Figure 3. *Relationship between tile resolutions and perspective view. (Source: SRI International)*

Id. at GOOG 349; *see also* '189 patent, Fig. 7 (perspective viewpoint). Perspective provides the illusion of depth, and is a technique for representing three-dimensional terrain on a flat surface. *See, e.g.,* Chang Decl., Ex. 22 at GOOG 350 (equating "perspective" views with "3-D" views). Thus, even without elevation data, image data can be used to represent three-dimensional terrain.

34. In its motion, Skyline also argues that there must be "detailed disclosure[s]" in the prior art to enable one of ordinary skill in the art to practice the invention of claims 1 and 12. *See, e.g.,* Skyline Validity Motion at 2. However, the '189 patent itself does not include such "detailed disclosure[s]," and if these disclosures are really necessary, then the patent would be invalid for lack of enablement. In particular, the '189 patent does not disclose any source code

for carrying out the invention, and it does not disclose how elevation data should be interactively downloaded separately from its combination with image data.

35. Skyline is both unfairly overestimating the complexity of the problem addressed by the ’189 patent and underestimating the person of ordinary skill in the art. First, the ’189 patent is not about *rendering* three-dimensional terrain (which is very complex and computationally intensive). Indeed, the ’189 patent acknowledges that “[c]omputer rendering of three-dimensional terrain images is well known in the art.” ’189 patent at col. 1:41-42. The ’189 patent is about *providing* the data blocks needed for rendering three-dimensional terrain to a renderer in an efficient way. These data blocks could have elevation data, they could have image data, they could have data for existing objects or structures, or they could have all three. In fact, each and every data block described in the preferred embodiments of the ’189 patent has both image data and elevation data. Furthermore, there is nothing different about *downloading* a data block with elevation data versus downloading a data block with image data. It is a data block, and exactly the same downloading functions and routines can be used for elevation and image data. Second, the solution claimed by the ’189 patent is to store a small number of data blocks in local memory, and then interactively download additional, higher resolution data blocks from a hierarchical data structure on a remote server as needed for rendering. There is nothing particularly difficult about implementing the downloading portion of this solution, and a person of ordinary skill in the art would have been more than capable of writing source code to achieve this end without any further direction. Indeed, given that there is no source code disclosed in the ’189 patent, and only basic flow charts, this appears to have also been the conclusion of the named inventors.

36. When the Court’s claim constructions are interpreting both correctly and

consistently and the correct standard for a person of ordinary skill in the art is applied, then the prior art invalidates at least claims 1 and 12 of the ’189 patent.

TERRAVISION

37. TerraVision was a system for visualization of terrain. It allowed a user to view and explore, in real time, a synthetic recreation of a real landscape created from elevation data and a large number of aerial or satellite images of that landscape.

38. There are at least three relevant prior art TerraVision references: (1) the public use of the TerraVision application; (2) the publication of the MAGIC Final Report in May 1996 (Chang Decl., Ex. 21); and (3) the publication of the MAGIC IEEE Article in May/June 1996 (*Id.*, Ex. 22).[1]

39. In my opinion, the public use of the TerraVision application anticipated at least claims 1 and 12 of the ’189 patent. *See, e.g.*, Feiner SJ Decl., Ex. D at ¶¶ 21-62, 85-87, 98-112. This opinion is not just based on my observation of TerraVision at SIGGRAPH ’95. Rather, it is supported by the testimony of Stephen Lau (one of the inventors of TerraVision, and the person who actually demonstrated TerraVision at SIGGRAPH ’95), by the source code for the TerraVision application, by a video illustrating TerraVision in operation, and by other documents corroborating the public use of TerraVision, not just at SIGGRAPH ’95, but at numerous other conferences and symposia. *Id; see also* Chang Decl., Exs. 14, 20-26, 30-34. Moreover, the TerraVision demonstration I observed at SIGGRAPH ‘95 was not a video, but an interactive demonstration in which the system was used to explore the database in real time.

40. Furthermore, in my opinion, the publication of the MAGIC Final Report in May

[1] The TerraVision Tech Note (Chang Decl., Ex. 23) is also a relevant TerraVision prior art reference, but for the purpose of claims 1 and 12, it is not necessary to also rely on this reference in light of the other references. Relevant disclosures regarding the TerraVision Tech Note are described in my expert report. Feiner SJ Decl., Ex. D.

1996 and the publication of the MAGIC IEEE Article in May/June 1996 also anticipate at least claims 1 and 12, at least to the extent those claims can be read so broadly as to capture Google Earth. *See, e.g.*, Feiner SJ Decl., Ex. D at ¶¶ 21-62, 85-87, 98-112.

41. Finally, in my opinion, the public use of the TerraVision application, the publication of the MAGIC Final Report and the publication of the MAGIC IEEE Article, either individually or in combination, at least render obvious claims 1 and 12. *See id.*

A. Anticipation

1. All Asserted Claims: "providing data blocks describing three-dimensional terrain"

42. The asserted claims recite a method of or apparatus for "providing data blocks describing three-dimensional terrain data." The Court has construed "data block describing three-dimensional terrain" to mean "a block or collection of data or digital information that represents or describes a section of three-dimensional terrain at a particular resolution level and that includes any additional data overlaid on the digital image of the terrain, such as altitude, labels or optional objects." Chang Decl., Ex. 4 at 9-12.

43. The TerraVision application, the MAGIC Final Report and the MAGIC IEEE Article all disclosed "providing data blocks describing three-dimensional terrain." TerraVision was a high-speed graphics application that allowed a user to interact in real time with a synthetic 3D photo-realistic view of a large terrain. This application accessed "tiles" representing both elevation and image data. Tiles with elevation data were referred to as digital elevation model ("DEM") tiles and tiles with image data were referred to as orthographic image ("OI") tiles. Both DEM tiles and OI tiles were data blocks describing three-dimensional terrain data as construed by the Court. They were data blocks or collections of digital information that represented or described a section of three-dimensional terrain, that is the "surface features of an

area of land, an object, or a material, including color, elevation, and existing objects or structures on the land, object or material." Moreover, both DEM tiles and OI tiles were provided, as rendering of the terrain on the screen in TerraVision was accomplished by combining the DEM and OI tiles for the selected area at the appropriate resolution, as illustrated in Figure 4 of the MAGIC IEEE Article:

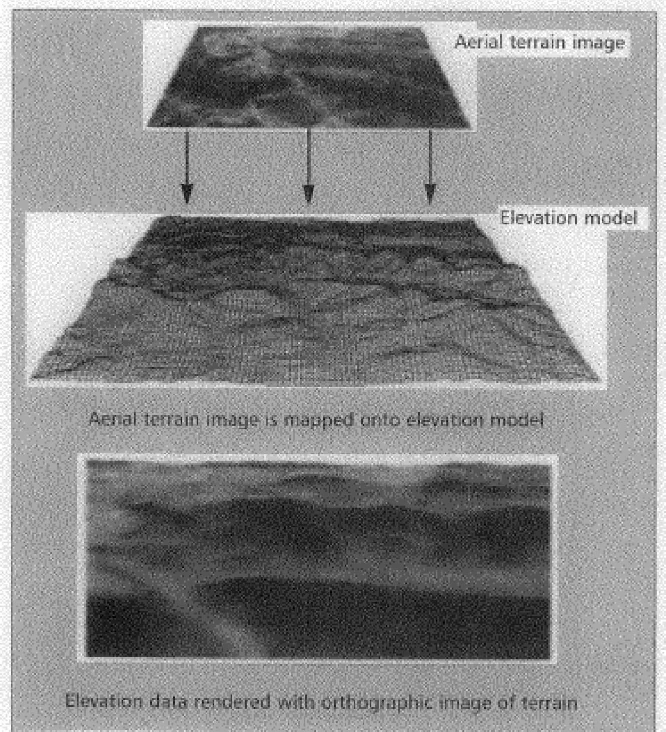

■ Figure 4. *Mapping an ortho-image onto its digital elevation model. (Source: SRI International)*

Chang Decl., Ex. 22 at GOOG 350.

2. All Asserted Claims: "renderer"

44. The asserted claims further indicate that the "data blocks describing three-dimensional terrain" are provided to a "renderer." The Court has construed the "renderer" to be a "software and/or hardware object that performs at least the following functions: (1) determining and providing to another object the required coordinates in the terrain along with a respective resolution level; (2) receiving the data blocks corresponding to the specified

coordinates; and (3) using the received data blocks to display a three-dimensional image." *Id.*, Ex. 4 at 26-32.

45. In my opinion, the TerraVision application disclosed a "renderer" as construed by the Court. The TerraVision source code confirms that the TerraVision application had software that performed the functions of (1) determining and providing to another object the required coordinates in the terrain along with a respective resolution level; (2) receiving the data blocks corresponding to the specified coordinates; and (3) using the received data blocks to display a three-dimensional image. Specifically, ThreeDWidgetGenerateVisible calls ThreeDWidgetCalcVisibility to create a quadtree identifying all the visible tiles within the frustum up to a resolution appropriate for the view matrix. Each tile is identified by an *x* coordinate, a *y* coordinate and a resolution level (e.g., QuadTile). ParseQuadTree (another object) receives the quadtree from ThreeDWidgetGenerateVisible (part of the renderer). ParseQuadTree then provides a list of the "leaf" tiles resident in memory to ThreeDWidgetGenerateVisible, which uses these tiles to create the underlying polygonal mesh that will be used to display the scene. Finally, the rendering thread will actually display the three-dimensional image based on the tiles provided.

46. Moreover, to the extent that Skyline seeks to interpret the term "renderer" more broadly than the Court's construction in an attempt to capture Google Earth, then in my opinion, the MAGIC Final Report and the MAGIC IEEE Article also disclosed a "renderer." In particular, Skyline applies a different interpretation of this term for the purposes of infringement and validity. If the Court accepts Skyline's interpretation for infringement, then in my opinion, the MAGIC Final Report and the MAGIC IEEE Article also disclosed a "renderer." I note that

regardless of what interpretation is applied, the TerraVision application had a renderer as discussed above.

47. The MAGIC Final Report and the MAGIC IEEE Article each disclose a search algorithm used to identify all the visible tiles at the appropriate resolution level. *See, e.g.*, Chang Decl., Ex. 21 at GOOG 363 (identifying tile visibility and rendering threads) & Ex. 22 at GOOG 350 ("A high-speed search algorithm is used to identify the tiles required to render a given view."). They further disclose that a small fraction of the available tiles are stored in local cache so that TerraVision is able to display a new view at any time, no matter how quickly the user moves. *Id.*, Ex. 21 at GOOG 364 (small fraction of tiles kept in local cache) & Ex. 22 at GOOG 351 (same). Based on these disclosures, a person of ordinary skill in the art would understand that the renderer must request and receive these data blocks from local memory and use them to display a three-dimensional image.

3. *All Asserted Claims: "data blocks belonging to a hierarchical structure"*

48. The asserted claims further recite "data blocks belonging to a hierarchical structure which includes blocks at a plurality of different resolution levels." The Court has construed "data blocks belonging to a hierarchical structure" as "data blocks that are organized into multiple levels of resolution, whereby each level contains data blocks at the same resolution, and each successive level contains data blocks of a higher resolution than those in the preceding level." Chang Decl., Ex. 4 at 12-15.

49. The TerraVision application, the MAGIC Final Report and the MAGIC IEEE Article all disclosed "data blocks belonging to a hierarchical structure which includes blocks at a plurality of different resolution levels," as that phrase has been construed by the Court. TerraVision used a multiresolution hierarchy or pyramid of increasingly lower-resolution

representations of the equal-sized tiles. *See, e.g., id.*, Ex. 21 at GOOG 359 & Ex. 22 at GOOG 349. Each level in the pyramid was at half the resolution of the previous level. *Id.* This is illustrated in Figure 3 of the MAGIC IEEE Article:

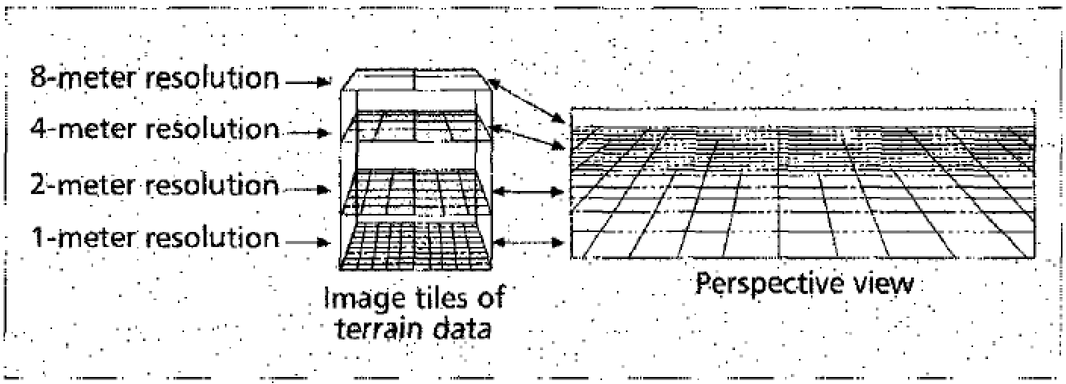

■ Figure 3. *Relationship between tile resolutions and perspective view. (Source: SRI International)*

Id., Ex. 22 at GOOG 349.

50. The TerraVision source code also confirms that TerraVision had "data blocks belonging to a hierarchical structure which includes blocks at a plurality of different resolution levels," as that phrase has been construed by the Court. The ThreeDWidgetCalcVisibility function in the TerraVision source code creates a quadtree identifying tiles at a plurality of different resolution levels. It can do this because the tiles on the remote servers have been organized as pyramids. Thus, the quadtree created by ThreeDWidgetCalcVisibility and the pyramids on the remote servers are organized into multiple levels of resolution, where each level has tiles at the same resolution, and each successive level has tiles of a higher resolution than those in the preceding level.

4. *All Asserted Claims: "receiving from the renderer one or more coordinates in the terrain along with indication of a respective resolution level"*

51. The Court has construed "coordinates in the terrain" as "a set of numerical values that identifies a particular location in the terrain," and "terrain" as "the surface features of an area of land, an object, or a material, including color, elevation, and existing objects or structures on the land, object or material." Chang Decl., Ex. 4 at 17-23. The Court has also clarified that "receiving from the renderer" means "something distinct from the renderer receiving from the renderer." *Id.*, Ex. 5 at 8-10.

52. The TerraVision application disclosed the step of "receiving from the renderer one or more coordinates in the terrain along with indication of a respective resolution level." As discussed above, ThreeDWidgetGenerateVisible calls ThreeDWidgetCalcVisibility to create a quadtree with all of the visible tiles within the frustum up to a resolution appropriate for the view matrix. Each tile is identified by an *x* coordinate, a *y* coordinate and a resolution level (QuadTile). ParseQuadTree (another object) receives this quadtree from ThreeDWidgetGenerateVisible.

53. The MAGIC Final Report and the MAGIC IEEE Article also disclosed this step. Tiles are identified by an *x* coordinate, a *y* coordinate and a resolution level. *See, e.g.*, Chang Decl., Ex. 21 at GOOG 359 & Ex. 22 at GOOG 349, 350 (noting that OI tiles and DEM tiles are "precisely" aligned "with a world coordinate system as well as with each other"). The (*x*,*y*) coordinates are a set of numerical values that identify the location in the terrain. *Id.* Thus, TerraVision references tiles using "coordinates in the terrain," as defined by the Court. A high-speed search algorithm is used to find all the visible tiles at the appropriate resolution level. *See, e.g.*, *id.*, Ex. 21 at GOOG 363 (identifying tile visibility and rendering threads) & Ex. 22 at

GOOG 350 ("A high-speed search algorithm is used to identify the tiles required to render a given view.").

5. *All Asserted Claims: "providing the renderer with a first data block which includes data corresponding to the one or more coordinates, from a local memory"*

54. The Court has construed "first data block" as "the first data block provided to the renderer from the local memory corresponding to the specified coordinates." *Id.*, Ex. 4 at 15-17. The Court has construed "local memory" as "memory easily accessible to the user's processor, either because it is physically part of the processor or is attached directly thereto, and distinct from the memory of the remote server from which data must be downloaded." *Id.* at 32-34. The Court has also clarified that "providing the renderer" means "something distinct from the renderer providing to the renderer." *Id.*, Ex. 5 at 8-10.

55. The TerraVision application disclosed "providing the renderer with a first data block which includes data corresponding to the one or more coordinates, from a local memory." ThreeDWidgetGenerateVisible calls ParseQuadTree, which identifies a "first data block which includes data corresponding to the one or more coordinates, from a local memory." Specifically, ParseQuadTree (something distinct from the renderer) identifies the visible tiles resident in memory and provides them to the renderer.

56. The MAGIC Final Report and the MAGIC IEEE Article also disclosed this step. TerraVision keeps a small fraction of the available tiles in a local cache (i.e., in local memory). *See, e.g.*, Chang Decl., Ex. 21 at GOOG 364 (small fraction of tiles kept in local cache) & Ex. 22 at GOOG 351 (same). Moreover, by using local memory (i.e., by providing the first data block to the renderer), TerraVision is able to display a new view at any time, no matter how quickly the user moves, although that view may not be at the desired resolution. *Id.* Both the visible tiles in memory and the tiles to be requested are identified using a "coarse-to-fine" strategy. *Id.* This

strategy manifests itself by the affected portion of the rendered image appearing "fuzzy" for a brief period of time before additional data blocks are downloaded. *Id.*

6. All Asserted Claims: "downloading from a remote server one or more additional data blocks at a resolution level higher than the resolution level of the first data block which include data corresponding to the one or more coordinates if the provided block from the local memory is not at the indicated resolution level"

57. The Court has construed "downloading" to mean "requesting over a network from a separate computer and receiving on a local computer." Chang Decl., Ex. 5 at 4-8. The Court has construed "downloading … if the provided data block from the local memory is not at the indicated resolution level" to mean "downloading … upon some determination that the block provided from local memory is not at the indicated resolution level." *Id.* at 10-12.

58. The TerraVision application disclosed "downloading from a remote server one or more additional data blocks at a resolution level higher than the resolution level of the first data block which include data corresponding to the one or more coordinates if the provided block from the local memory is not at the indicated resolution level," as construed by the Court. ThreeDWidgetGenerateRequests calls ThreeDWidgetCalcVisibility to generate a quadtree of requested tiles at up to the appropriate resolution based on a "bloated" view matrix (i.e., a view matrix specifying an expanded version of the view frustum used by ThreeDWidgetGenerateVisible). This quadtree is passed to GenerateAndSendRequests as a sorted list, in coarse-to-fine resolution order. The tiles in the list include those in local memory, as well as those on a remote server. GenerateAndSendRequests then loops through the list, in coarse-to-fine order, to determine whether or not each tile is in local memory. If a tile in the list is already in local memory, then GenerateAndSendRequests will cause the usage time on the tile to be updated, so as to prevent the tile from being deleted from local memory, and it will not attempt to download that tile from a remote server. When GenerateAndSendRequests

encounters a tile that is not in local memory (i.e., as it continues its coarse-to-fine traversal), then it will place that tile on a download request list. The download request list is processed by the tile set manager, which actually downloads the tiles from a remote server. The tile set manager supports several alternative methods for downloading tiles, including the use of ISS protocols and the use of HTTP to access tiles addressed as URLs. GenerateAndSendRequests will request that additional tiles at a resolution level higher than the resolution level of the first tile provided from local memory be downloaded *if* the provided tile from local memory is not at the indicated resolution level. TerraVision thus downloaded additional tiles from the remote server "upon some determination that the block provided from local memory is not at the indicated resolution level."

59. In the TerraVision application, both OI tiles and DEM tiles were requested and received from the remote server. DEM tiles were downloaded from the remote server when the user selected a data set, and then stored in local cache memory since there was only a relatively small number of DEM tiles. In fact, DEM tiles were requested in coarse-to-fine order in TsRequestDems, and the download list was created and processed using the exact same tsmReqTile and tsmStopReqTiles functions used for OI tiles. OI tiles were not downloaded when the user selected an initial data set, but would rather be downloaded only if the provided block in local memory was not at the indicated resolution level. OI tiles (like DEM tiles) were "data blocks [describing three-dimensional terrain]," and the downloading of these additional OI tiles satisfies the claim. The OI tiles described the appearance of the terrain and were needed (along with the DEM tiles) to render the three-dimensional terrain in the TerraVision application. *See also* ¶¶ 42-43 above.

60. In contrast, a person of ordinary skill in the art would understand from either the MAGIC Final Report or the MAGIC IEEE Article that *both* OI tiles and DEM tiles were "streamed" from the remote server interactively. These references disclosed that all data—including elevation data—was stored remotely. *See* Chang Decl., Ex. 21 at GOOG 362 ("[TerraVision] combines elevation data, aerial photographs, models of buildings and models of vehicles whose positions were obtained using GPS receivers, all stored in a remote terrain database (ISS) accessed via a high-speed network.") & Ex. 22 at GOOG 350 ("The ISS stores, organizes, and retrieves the processed imagery and elevation data required by TerraVision for interactive rendering of the terrain."). Moreover, there was only one method disclosed in these references for downloading data from the remote server, i.e., data was downloaded on an as-needed basis. *See id; see also id.*, Ex. 21 at GOOG 363-64 (describing coarse-to-fine downloading and rendering strategy) & Ex. 22 at GOOG 351 (same). There was nothing in either the MAGIC Final Report or the MAGIC IEEE Article indicating that all DEM tiles should be downloaded initially or stored locally. Rather, these references repeatedly refer to downloading "tiles," which one of ordinary skill would understand to refer to both OI and DEM tiles. Consequently, one of ordinary skill in the art would be able to download both OI and DEM tiles in the manner and order described in the asserted claims. Indeed, Skyline does not appear to contest that these references sufficiently disclose how to download OI tiles in the manner and order described in the asserted claims. In addition, there is nothing about the way in which elevation data is stored on a server in TerraVision or in the preferred embodiments of the '189 patent that would make such downloading any different than downloading color data.

61. Furthermore, the MAGIC Final Report specifically disclosed a "determination that the block provided from local memory is not at the indicated resolution level." In particular,

the MAGIC Final Report states that "[t]here is a 'tile prediction' thread for predicting the user's movement and determining the set of tiles that will be visible in the predicted viewpoint, and a 'tile requesting' thread for requesting those tiles from the ISS (eliminating those that are already in memory)." Chang Decl., Ex. 21 at GOOG 364. This step of "eliminating those that are already in memory" is a reference to the determination made by TerraVision of whether the block provided from local memory is at the indicated resolution level—TerraVision would download additional data blocks if the block in memory was not at the indicated resolution level. In addition, the MAGIC IEEE Article states that imagery and elevation data were stored remotely for "interactive rendering of the terrain." *See id.*, Ex. 22 at GOOG 350. It also disclosed rendering the highest resolution terrain possible while downloading additional, higher resolution tiles. *Id.* at GOOG 351.

B. Obviousness

62. Even assuming that the downloading of additional OI tiles does not satisfy the Court's construction and that a person of ordinary skill in the art would be incapable of carrying out the directions in the MAGIC Final Report and the MAGIC IEEE Article to download both additional OI tiles and additional DEM tiles, in my opinion, claims 1 and 12 would have at least been obvious.

63. A person of ordinary skill in the art would have been motivated to combine the TerraVision application with the MAGIC Final Report and the MAGIC IEEE Article, since all these references refer to the same project and system. Thus, to the extent that Skyline claims that the MAGIC Final Report and MAGIC IEEE Article lack necessary "detailed disclosures" regarding a "renderer," these detailed disclosures are set forth in the TerraVision application and are thus fully enabled.

64. Moreover, both the MAGIC Final Report and the MAGIC IEEE Article expressly teach that both elevation and image data should be downloaded interactively, as needed by the user. They do not teach that all DEM tiles should be stored in local memory or that all DEM tiles should be downloaded for an area of interest prior to OI tiles. Thus, they do not "teach away" from the '189 patent.

65. The TerraVision application did download all DEM tiles for an area of interest prior to OI tiles. However, implementing the same coarse-to-fine strategy for downloading DEM tiles as used for downloading OI tiles would have involved minor modifications of the TerraVision source code, which already includes relevant functions and variables for this purpose. Whether these modifications would have affected the quality with which the scene would be rendered is not at issue—the '189 patent is concerned with providing the data blocks to the renderer, not with the quality of the rendering. The TerraVision application would have still been able to render a three-dimensional scene with these modifications.

66. Skyline's suggestion that the inventors of the TerraVision application could not solve this downloading problem is unsupported. Skyline first relies on outdated reports and legacy source code functions to allegedly show that in the TerraVision application, DEM tiles were not downloaded at all, but were stored locally. That was certainly true in the early stages of development for TerraVision, but the source code shows that at least as of December 14, 1994, there was a function (TsRequestDems) in TerraVision for downloading DEM tiles from a remote server. Chang Decl., Exs. 20, 34; *see also* Haight Decl., Ex. 24 (April 1995 Quarterly Rpt. at GOOG 26479 (covering period from Jan. 1, 1995 through Mar. 31, 1995)). Thus, to the extent that the TerraVision application used a "hack until we can get the TSM to send us the DEMs," this hack was superseded in December 1994 by source code that used the TSM to send DEM

tiles with the same mechanisms used to send OI tiles.

67. Furthermore, at least as of 1995, the default approach employed in the TerraVision application was to download DEM tiles from a remote server, not to obtain them locally. Chang Decl., Exs. 20, 34. As noted in TerraVisionInitDataSet, a special flag had to be set explicitly upon invocation of TerraVision in order to obtain DEM tiles locally (the "cheesy local DEMs") rather than to obtain them remotely by default from the server. *Id.* Thus, obtaining just the DEM tiles locally was not part of the regular operation of the application, at least as of 1995.

68. Initially downloading all of the DEM tiles for a small area of the terrain does not mean that the inventors of the TerraVision application were unable to download these tiles the same way as they downloaded the OI tiles. Rather, this "greatly simplifie[d] the resolution determination algorithm at a comparatively small cost, since the number of DEM tiles is relatively small." Haight Decl., Ex. 24 (April 1995 Quarterly Rpt. at GOOG 26479). The claims of the '189 patent, however, are not limited to any particular method for determining a resolution level—they rather require that the renderer specify "an indication of a respective resolution level."

69. Finally, it would have been obvious to a person of ordinary skill in the art familiar with the TerraVision application to also download the DEM tiles in the same manner as OI tiles were downloaded for a larger dataset (i.e., one with a larger amount of data). The reason the developers of TerraVision were able to download all of the DEM tiles initially for an area of interest was because "the number of DEM tiles is relatively small." *Id.* This would not necessarily be true of a larger dataset, and accordingly a person of ordinary skill would have had motivation to implement the modifications discussed above. It also would have been obvious to

use a larger data set. *See, e.g.,* Chang Decl., Ex. 21 ("Thus, a user can roam over arbitrarily large databases without having to wait for the entire database to be downloaded first.").

T_VISION

70. T_Vision was a method and device for the pictorial representation of space-related data, for example, geographical data of the earth. *See, e.g.*, Haight Decl., Ex. 51 (U.S. Patent No. 6,100,897 at Abstract).

71. There are at least three relevant prior art T_Vision references: (1) the Mayer patent, claiming priority from a German patent application filed in December 1995 (*Id.*, Ex. 51); (2) the publication of materials describing the T_Vision Project on the SIGGRAPH '95 Multimedia CD-ROM in July 1995 (Mewes Decl., Exs.2 & 3; Haght Decl. Ex.30); and (3) the public use of T_Vision at SIGGRAPH '95.

72. In my opinion, the Mayer patent and the T_Vision Project materials on the SIGGRAPH '95 Multimedia CD-ROM anticipate at least claims 1 and 12 of the '189 patent, at least to the extent those claims can be read so broadly as to capture Google Earth. *See, e.g.,* Feiner SJ Decl., Ex. D at ¶¶ 113-131, 146-48, 159-68.

73. Furthermore, in my opinion, the public use of the T_Vision application at SIGGRAPH '95 also anticipates at least claims 1 and 12, at least to the extent those claims can be read so broadly as to capture Google Earth. *See id.* This opinion is not just based on my observation of T_Vision at SIGGRAPH '95, but is also based on the T_Vision Project materials, on the Mayer patent, on the Terra1995 a video illustrating T_Vision in operation, and on other documents corroborating the public use of T_Vision at SIGGRAPH '95. Haight Decl., Ex. 51; Mewes Decl., Exs. 2-4. In particular, the T_Vision project materials were distributed in connection with the public demonstration of the T_Vision application at SIGGRAPH '95 and

describe that system. The Mayer patent is based on a German patent application filed in December 1995, four months after SIGGRAPH ’95 (held in August 1995). It names as inventors the same individuals identified in the T_Vision Project materials. It includes some of the same figures as the T_Vision Project materials. And the disclosures in the patent are consistent with the disclosures in the T_Vision Project materials (albeit, the patent provides much greater detail and also discloses alternative embodiments). The T_Vision video is dated 1995, the same year as SIGGRAPH ’95, and also includes some of the very same footage about T_Vision included in the T_Vision SIGGRAPH ’95 Project materials. Finally, the T_Vision demonstration I observed at SIGGRAPH ‘95 was not a video, but an interactive demonstration in which the system was used to explore the database in real time.

74. It is also clear that the T_Vision application was not a “mere concept” when it was demonstrated at SIGGRAPH ’95. The T_Vision project materials clearly describe “The (already existing and working) Prototype.” Mewes Decl. Ex. 3 (RENDERER.HTM at p. 1). A working prototype is more than a concept. These materials further describe the “Task” as developing “a renderer which visualize[s] a worldwide distributed database with unlimited geometry and textures in realtime.” *Id.* After discussing the working prototype, the developers do identify some problems in completely fulfilling this “Task,” but none have any relevance to the claims at issue. *Id.* at p. 2. In particular, there is no indication at all that the developers of T_Vision had any problem whatsoever in downloading or rendering elevation data. *Id.*

75. Moreover, the statement from one of the developers of T_Vision that he intended to “throw away all the code I have written for this prototype and replace it with a real object oriented full feature distributed high performance real time database” did not mean that T_Vision was just a concept. *See* Mewes Decl., Ex. 3 (TERRABAS.HTM at p. 3). As an initial matter,

this statement was made about the *database* alone, not the entire program. *See id*. But further, this shows that the developers of T_Vision did have a working prototype database. *Id.* (TERRABAS.HTM at pp. 1-2). That they thought they could build a better database does not mean that the working prototype database, let alone T_Vision, was merely a "concept."

76. Finally, in my opinion, the Mayer patent, the publication of the T_Vision Project materials, and the public use of the T_Vision application, either individually or in combination, at least render obvious claims 1 and 12 of the '189 patent. *See* Feiner SJ Decl., Ex. D at ¶¶ 113-131, 146-48, 159-68.

A. Anticipation

1. All Asserted Claims: "providing data blocks describing three-dimensional terrain"

77. The asserted claims recite a method of or apparatus for "providing data blocks describing three-dimensional terrain data." The Court has construed "data block describing three-dimensional terrain" to mean "a block or collection of data or digital information that represents or describes a section of three-dimensional terrain at a particular resolution level and that includes any additional data overlaid on the digital image of the terrain, such as altitude, labels or optional objects." Chang Decl., Ex. 4 at 9-12.

78. The T_Vision Project materials disclosed a "method of providing data blocks describing three-dimensional terrain to a renderer." These materials state that the T_Vision application provided a "virtual globe" that was "modeled from high resolution spatial data and textured with high resolution satellite images." Mewes Decl., Ex. 3 (TVISION.HTL at p. 1). The "high resolution spatial data" is elevation data and it is "textured" or overlaid with "high resolution satellite images." *Id.* Moreover, these materials further clarify that both image data and elevation data were used to render the terrain: they describe a database with "pairs of index

and data files containing 128x128 pixel texture images (surface, clouds) and 16x16 point elevation data." *Id.* (TERRABAS.HTM at p. 1); *see also id.* ("Currently the source data consists of around 10 GB of image and DEM data, covering the whole world in 4 km/pixel, USA and Europe in 1km, Japan in 50m, some areas in USA and Germany in 50m, and parts of Berlin and Tokyo down to 30cm."). The rendering of three-dimensional terrain is also illustrated in the T_Vision Project materials:

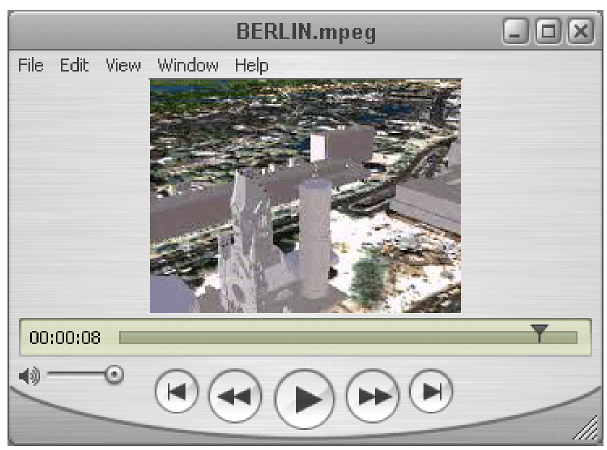

Mewes Decl., Ex. 3 (TVISION.HTL at p. 1); *see also id.*, Ex. 2 (TERRA_S.mpeg & BERLIN.mpeg).

79. That the T_Vision application publicly demonstrated at SIGGRAPH ’95 rendered three-dimensional terrain data was also graphically demonstrated and described in the T_Vision Terra1995 video:

> By switching the surface off, we can observe this process more easily. As the distance between us and the Earth increases, the high resolution data is removed from the memory and is replaced with new data for the wider field of view. Out of these different levels of altitude data, we compute the tectonic surface of the Earth and then project the corresponding satellite images onto it. An asynchronous and anticipatory loading strategy always guarantees a steady frames per second.

Mewes Decl., Exs. 4-5. In addition to the regular textured view, this portion of the video shows both wireframe and shaded, untextured views of the terrain that clearly indicate that it was rendered from a 3D model whose elevation varies across the terrain:

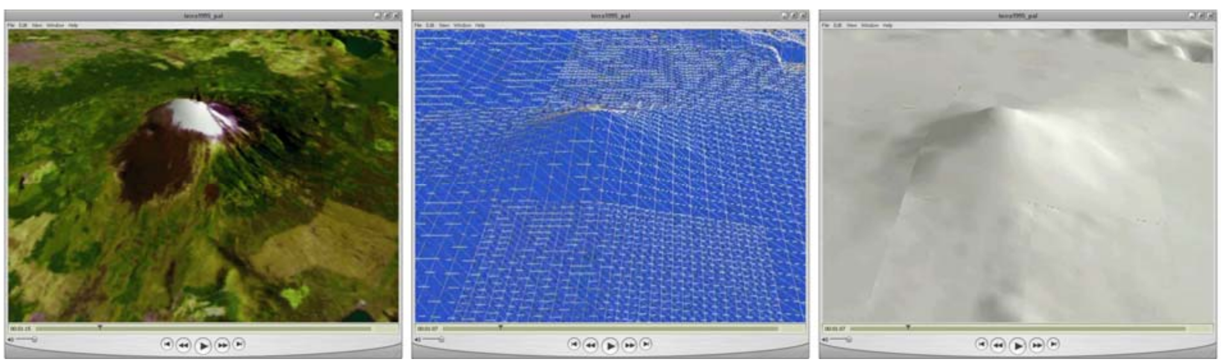

80. The Mayer patent also disclosed a “method of providing data blocks describing three-dimensional terrain to a renderer.” The patent is described as, “[a] method and device for the pictorial representation of space-related data, for example, geographical data of the earth.” Haight Decl., Ex. 51 (’897 patent at Abstract). Moreover, the “representation” referred to may be “based on a three-dimensional geometrical model.” *Id.* at col. 8:56-57. Furthermore, the Mayer patent specifically states that elevation data, as well as image data, is used for certain pictorial representations: “In order to show the field of view with this image resolution a height value is required every 150 m and an image value of a surface every 15 m.” *Id.* at col. 8:14-17.

It also repeatedly refers to the visualization of topographic information, including the overlay of color data on a "topographical grid network of the earth surface." *See, e.g.*, *id.* at col. 1:7-10 & 9:18-43. "In the topographical grid model the polygon grid imitates the topography of the surface" (and thus represents elevation data). *Id.* at col. 4:36-38. The topographical grid model is described as "two-dimensional," but this is a reference to how the data is stored (as a two-dimensional array of values), not to the type of data rendered. A person of ordinary skill in the art would recognize from these disclosures that the techniques disclosed in the Mayer patent were intended to be used to display three-dimensional terrain, in addition to other objects.

2. All Asserted Claims: "renderer"

81. The asserted claims further indicate that the "data blocks describing three-dimensional terrain" are provided to a "renderer." The Court has construed the "renderer" to be a "software and/or hardware object that performs at least the following functions: (1) determining and providing to another object the required coordinates in the terrain along with a respective resolution level; (2) receiving the data blocks corresponding to the specified coordinates; and (3) using the received data blocks to display a three-dimensional image." *Id*., Ex. 4 at 26-32.

82. To the extent that Skyline seeks to interpret the term "renderer" more broadly than the Court's construction in an attempt to capture Google Earth, then in my opinion, the T_Vision Project materials, the T_Vision application and the Mayer patent also disclosed a "renderer."

83. The T_Vision Project materials describe data blocks with a Global Area Identifier corresponding to "coordinates in the terrain along with a respective resolution level." Mewes Decl., Ex. 3 (TERRABAS.HTM at pp. 1-2). Each "tile" or "patch" has "bounds generated by

binary subdivision of the whole coordinate system." *Id.* at p. 1. "The GAI can be seen as a kind of telephone number for reaching a particular sector of the planet. The number of digits corresponds to the level of detail; the higher the number, the finer the resolution." *Id.* This same system is disclosed in the Mayer patent. *Compare* Haight Decl., Ex. 51 ('897 patent at col. 8:28-67 & Figs. 4-6).

84. The T_Vision Project materials further state that "[t]he renderer computes the GAIs according to the field of view and makes a simple query." That is, it provides the coordinates in the terrain and an indicated resolution level to another object. Mewes Decl., Ex. 3 (TERRABAS.HTM at p. 2); *see also id.* (RENDERER.HTM at p. 1) ("Then it requests the data for a special location with an appropriate resolution."). The renderer also receives data blocks from local memory and renders these data blocks for view. *Id.* (RENDERER.HTM at pp. 1-2).

85. Likewise, the Mayer patent discloses a "node 3" which is the client computer (an SGI Onyx in the preferred embodiment). Part of this "node 3" computer provides the coordinates in the terrain and an indicated resolution level to another object. Haight Decl., Ex. 51 ('897 patent at col. 8:28-38). For geographically related data, "[t]he representation may in this case be carried out both according to cartographic points of view or also as a globe." *Id.* at col. 4:14-17. Moreover, "[i]f the node 3 then ascertains that the required screen resolution has not been achieved with the centrally stored data [i.e., the data in the local memory of node 3], it divides the field of view according to the model of the quadrant tree into four sections and checks each section to see whether, by representation of the data contained in the sections, the required image resolution has been achieved." *Id.* at col. 7:45-50; *see also id.* at col. 2:11-22 (first data set centrally stored and the field of view is shown) & col. 3:27-33 ("After each transmission and central storage of data, an image representation results, even if the data are

insufficient to make possible the desired image resolution. As a result, even if the method is interrupted due to an alteration in the field of view and newly begun for a new field of view, the data for an image, even at low resolution, are always available.”).

3. *All Asserted Claims: “data blocks belonging to a hierarchical structure”*

86. The asserted claims further recite “data blocks belonging to a hierarchical structure which includes blocks at a plurality of different resolution levels.” The Court has construed “data blocks belonging to a hierarchical structure” as “data blocks that are organized into multiple levels of resolution, whereby each level contains data blocks at the same resolution, and each successive level contains data blocks of a higher resolution than those in the preceding level.” Chang Decl., Ex. 4 at 12-15.

87. The T_Vision application has “data blocks belonging to a hierarchical structure which includes blocks at a plurality of different resolution levels,” as that phrase has been construed by the Court. As disclosed in the T_Vision Project materials, the T_Vision application used a “multi-layered database” that was “organized as a quadtree, containing higher levels of detail as you descend down the tree.” Mewes Decl., Ex. 3 (RENDERER.HTM at p. 1). Levels of detail refer to resolution levels.

88. Likewise, the Mayer patent explains that the “sub-division of the image into sections with different spatial resolutions is preferably effected according to the method of a binary or quadrant tree.” Haight Decl., Ex. 51 (’897 patent at Abstract). This is illustrated in Figures 4–6 of the Mayer patent, which show data blocks belonging to a hierarchical structure. *Id.* at Figs. 4–6. For example, as explained in the Mayer patent: “FIG. 6 shows a sub-division according to an octant tree for a representation based on a three-dimensional geometrical model. Here a section 14 o[f] a space is sub-divided into eight spatial sub-sections 15. By means of the

method according to the invention, consequently here also the data of just the spatial areas are called up in a higher accuracy, at which it is required in order to obtain the desired image resolution." *Id.* at col. 8:55-62.

89. At a minimum, these disclosures satisfy the "hierarchical structure" limitation as that limitation has been interpreted by Skyline in an attempt to capture Google Earth. Moreover, there is nothing in these disclosures indicating that anything but a strict hierarchy was used.

4. *All Asserted Claims: "receiving from the renderer one or more coordinates in the terrain along with indication of a respective resolution level"*

90. The Court has construed "coordinates in the terrain" as "a set of numerical values that identifies a particular location in the terrain," and "terrain" as "the surface features of an area of land, an object, or a material, including color, elevation, and existing objects or structures on the land, object or material." Chang Decl., Ex. 4 at 17-23. The Court has also clarified that "receiving from the renderer" means "something distinct from the renderer receiving from the renderer." *Id*., Ex. 5 at 8-10.

91. As discussed above, the T_Vision Project materials expressly disclosed that "[t]he renderer computes the GAIs according to the field of view and makes a simple query." Mewes Decl., Ex. 3 (TERRABAS.HTM at p. 2). Thus, the "coordinates in the terrain along with indication of respective resolution level" (i.e., GAIs) are sent by the renderer, and (by implication), received by another object.

92. The Mayer patent states that part of the "node 3" computer requests "sections" of the field of view by providing coordinates along with an indication of a respective resolution level. As described in the Mayer patent: "Fig. 4 shows the formation of an address of a section using the model of a quadrant tree for sub-division of the field of view 11. In the first sub-division of the field of view 11 into four sections 12, these are identified clockwise with the

numerals 0 to 3. If a section is further sub-divided, the individual sub-sections 13 are numbered in the same way and the numbers thus obtained are prefixed to the numbers of the master section. With a permanently identical resolution of for example 128x128 points per section, the number of points of the section number is at the same time an indication of the level of spatial precision of the data." *Id.* ('897 patent at col. 8:28-38).

5. All Asserted Claims: "providing the renderer with a first data block which includes data corresponding to the one or more coordinates, from a local memory"

93. The Court has construed "first data block" as "the first data block provided to the renderer from the local memory corresponding to the specified coordinates." Chang Decl., Ex. 4 at 15-17. The Court has construed "local memory" as "memory easily accessible to the user's processor, either because it is physically part of the processor or is attached directly thereto, and distinct from the memory of the remote server from which data must be downloaded." *Id*. at 32-34. The Court has also clarified that "providing the renderer" means "something distinct from the renderer providing to the renderer." *Id*., Ex. 5 at 8-10.

94. The T_Vision Project materials provide that the renderer "requests the data for a special location with an appropriate resolution…. If you approach too fast you will get a coarse image, but the frame rate is not affected." Mewes Decl., Ex. 3 (RENDERER.HTM at p. 1). A person of ordinary skill in the art would understand from this disclosure that the "coarse" image should be rendered from a first data block in local memory because the desired data block has not yet been downloaded. *See also id.* (TVISION.HTL at p. 2).

95. The Mayer patent states: "After each transmission and central storage of data, an image representation results, even if the data are insufficient to make possible the desired image resolution. As a result, even if the method is interrupted due to an alteration in the field of view and newly begun for a new field of view, the data for an image, even at low resolution, are

always available." Haight Decl., Ex. 51 ('897 patent at col. 3:27-33). The Mayer patent even refers to this coarse representation as a "first data set." *Id.* at col. 2:14-17.

96. Skyline asserts that the Mayer patent does not disclose a "local memory" that is easily accessible to the "processor." However, Figure 2 of the patent clearly shows a node 3 computer directly connected to an input medium (10) and at least one display:

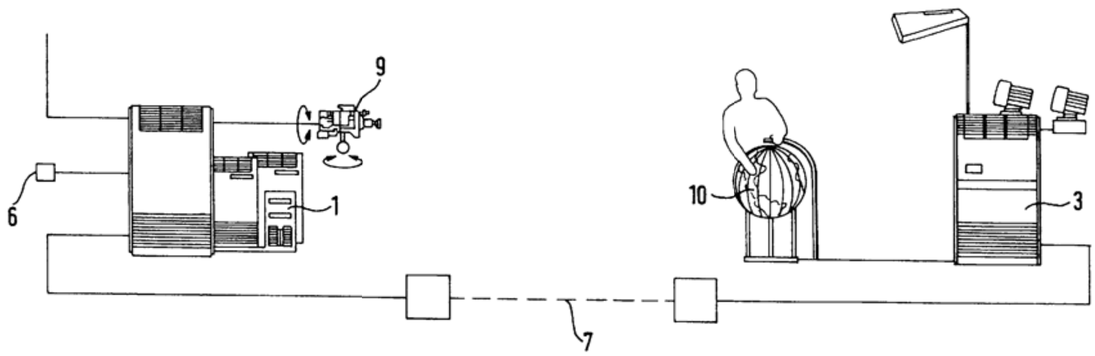

FIG. 2

See Haight Decl., Ex. 51 ('897 patent at col. 5:61-66, 6:64-7:1, 7:34-42 & Fig. 2). As shown in Figure 2, the user selects the field of view using a three-dimensional track ball (the input medium), connected to node 3. *See id.* at col. 7:3-7 & Fig. 2. Node 3 then determines if the data blocks needed to render that view are in its "central storage" (i.e., the local memory of node 3); if they are not, then node 3 calls up additional data blocks from the spatially distributed database (accessed through the collector network 6) via a remote server (node 1). *See, e.g., id.* at col. 3:27-33, 3:44-50, 7:45-59. A communications link (7) connects node 3 and node 1. *Id.*

97. The Mayer patent does disclose embodiments with one or more display units attached via a supply network 8. Skyline characterizes these display units as individual client computers. However, at least in the embodiment shown in Figure 2, it is clear that the user is not connecting to node 3 through some other computer—the three-dimensional track ball is directly connected to node 3, which is also shown as having a display. Indeed, this appears to be the

same set-up illustrated in the T_Vision Project materials:

Mewes Decl., Ex. 2 (TRACKER.mpeg (showing two displays, both with the same representation)). Thus, even assuming that the display units disclosed in the Mayer patent are separate computers, a person of ordinary skill in the art would have understood from this patent that the node 3 computer, with its central storage and processor, could be set up with its own display and directly connected to an input medium. In fact, the Mayer patent states that "systems of the company Silicon Graphics (SGI Onyx) were used as a node computer. This computer is capable of displaying more than 5,000,000 texturised triangles per second and consequently is suitable for rapid picture build-up." Haight Decl., Ex. 51 ('897 patent at col. 6:22-25). Moreover, Skyline's suggestion that the large track ball in Figure 2 is actually "remote" from node 3 is not accurate. The three-dimensional track ball—like the mouse on your computer—is simply an input medium.

6. All Asserted Claims: "downloading from a remote server one or more additional data blocks at a resolution level higher than the resolution level of the first data block which include data corresponding to the one or more coordinates if the provided block from the local memory is not at the indicated resolution level"

98. The Court has construed "downloading" to mean "requesting over a network from a separate computer and receiving on a local computer." Chang Decl., Ex. 5 at 4-8. The Court

has construed "downloading … if the provided data block from the local memory is not at the indicated resolution level" to mean "downloading … upon some determination that the block provided from local memory is not at the indicated resolution level." *Id.* at 10-12.

99. The T_Vision Project materials disclosed an application that downloaded additional, higher resolution data blocks as needed. This application was described as a "real-time rendering system" where remote data was "integrated unobtrusively into the user's system on the fly." Mewes Decl., Ex. 3 (TVISION.HTL at pp. 1-2). It displayed a coarse view if you moved too fast, since it had to download additional, higher resolution data blocks in order to display the view at the requested resolution level. *Id.* (RENDERER.HTM at p. 1).

100. There is nothing in these materials to indicate that only image data was downloaded. These materials explicitly state that the T_Vision database included data files with both image and elevation data, and accordingly, made clear that both types of data needed to be downloaded. *Id.* (TERRABAS.HTM at p. 1). As the only method disclosed for downloading data was to download it in "real-time" as needed by the user, a person of ordinary skill in the art would have understood that both image data and elevation data were downloaded in this manner.

101. The T_Vision Project materials include the following statement: "The database basically consists of pairs of index and data files containing 128x128 pixel texture images (surface, clouds) and 16x16 point elevation data. Geometry and Billboards are not stored at other places in the current implementation, but this will change in the future." I understand that Skyline asserts that this last sentence means elevation data was stored locally. This claim, however, is not supported. The prior sentence makes clear that the database already included elevation data: "The database basically consists of pairs of index and data files containing 128x128 pixel texture images (surface, clouds) *and 16x16 point elevation data*." *Id.* (emphasis

added). It was other "geometry" data (such as CAD models of buildings) and billboards that were "not stored at other places in the current implementation." *Id; see also id.* (RENDERER.HTM at p. 1 (database contains, "surface data (satellite imagery and aerial photographs), elevation data, transparent clouds, CAD-Models of buildings and Information billboards displaying names and current temperatures of selected cities")). Moreover, I disagree that this statement necessarily means even that the CAD models and billboards were stored locally. The T_Vision system establishes an ATM connection to "the server" that "provides the most up-to-date and highest resolution data required for the current field of view." *Id.* (TVISION.HTL). Thus, the CAD models and billboards may have been stored on this remote server instead of being collected from the spatially distributed data sources ("other places"). In that case, the geometry and billboards would still have needed to be downloaded to the client. Finally, even if one read this statement to mean that elevation data was stored locally, the T_Vision Project materials made clear that it could also be stored remotely and downloaded in real-time: "this will change in the future." *Id.* (TERRABAS.HTM at p. 1).

102. The T_Vision Project materials also disclosed that remotely stored data was accessed via NFS on an ATM-network. *Id*. at p. 2. Thus, regardless of whether or not the T_Vision application publicly demonstrated at SIGGRAPH '95 was connected to a network, these published materials disclosed a communications link. Moreover, there is also evidence indicating that the T_Vision application was connected to a network at SIGGRAPH '95, including testimony from Stephen Lau and documents showing that T_Vision had an ATM network connection set up at SIGGRAPH '95. *Id.*, Exs. 2, 6 (Lau Depo at 44:20-47:25, 84:5-16, 202:12-205:12) & 9 (Clinger Affidavit).

103. The Mayer patent also disclosed downloading additional, higher resolution data

blocks in real-time as needed by the user. For example, the patent states that "[i]f the resolution of the representation is below the desired image resolution, the field of view is divided into sections and an investigation is undertaken for each individual section to see whether the data within the section are sufficient for a representation with the desired image resolution. If this is not the case for one of the sections, *further data with a finer resolution are called up*, transmitted and centrally stored from at least one of the spatially distributed data sources, and the section is shown with the new data. In turn an investigation is carried out into sufficient image resolution and possibly a further sub-division of the tested section is carried out into further partial sections as described above." Haight Decl., Ex. 51 ('897 patent at col. 2:17-29); *see also id.* at col. 7:45-59. Moreover, as with the T_Vision Project materials, there is nothing to indicate that elevation data was treated any differently than image data.

104. The Mayer patent also explicitly states that a determination is made as to whether the block provided from local memory is not at the indicated resolution level before additional data blocks are downloaded—it tests whether the data within a section is sufficient for a representation with the desired image resolution. *See id.* If not, further data is downloaded. *Id.* Thus, the Mayer patent discloses a link between the condition of the provided block from local memory not being at the indicated resolution level and downloading additional data blocks. Moreover, "coarse" data blocks are downloaded first, and then progressively higher resolution data blocks are downloaded.

B. Obviousness

105. The Mayer patent, the T_Vision Project materials and the T_Vision application at minimum render obvious claims 1 and 12 of the '189 patent.

106. A person of ordinary skill in the art would have been motivated to combine the Mayer patent, the T_Vision Project materials and the T_Vision application, since all these references refer to the same project and system. Thus, to the extent that Skyline claims that the T_Vision Project materials or even the Mayer patent lacked necessary "detailed disclosures" regarding a "renderer" or a "hierarchical structure," these features were implemented in the T_Vision application, and thus were fully enabled.

107. Moreover, even assuming that elevation data was not downloaded in any of these references (though all of them indicate that it was), these references explicitly teach that elevation data, like image data, should be downloaded interactively in real time. Each of these references also includes sufficient disclosures that a person of ordinary skill in the art would have been enabled to download elevation data as well as image data—elevation data, like image data, is simply contained in data blocks, and the same processes can be used to download both types of data.

MIGDAL & COSMAN

108. The Migdal patent was filed on November 6, 1995, and issued on June 2, 1998. Haight Decl., Ex. 52. It describes a method and system for providing texture using a selected portion of a texture map. The Cosman article was presented at the IMAGE VII Conference in June 1994, and relates to global terrain texture. *Id.*, Ex. 28.

109. The Migdal patent was cited by the Examiner during prosecution of the '189 patent. In particular, the Examiner described Migdal as "a system and method for modeling 3D objects," and found that it disclosed:

> [A] method of providing data blocks (LOD generation block 1050, FIG. 10), describing three-dimensional terrain to a renderer (raster subsystem 224, FIG. 2), the data blocks belonging to a hierarchical structure which includes blocks at a plurality of different resolution levels (col. 9, ll.5-17), the method comprising:

- receiving from the renderer one or more coordinates in the terrain along with indication of a respective resolution level (col.16, ll.1-21);
- providing the renderer with a first data block which includes data corresponding to the one or more coordinates from a local memory (col.9, ll.5-14);
- downloading from a remote server one or more additional data blocks which include data corresponding to the one or more coordinates if the provided block from the local memory is not at the indicated resolution level (col.8, l.66-col.9, l.36 and FIG. 2. Local memory: texture memory 226. Remote server: graphics display system 200).

Chang Decl., Ex. 2 at GOOG 105 (7/5/01 Office Action). The Examiner noted that Migdal did not disclose "downloading a block at a resolution level higher than the resolution level of the first block," because Migdal "provides the data block with the highest resolution first." *Id.* at GOOG 108-09. However, the Examiner found that it would have been obvious to reverse the order of elements. *Id.* Accordingly, he rejected application claims 3 and 33 (corresponding to claims 1 and 12 of the '189 patent) as obvious.

110. In response, Skyline made no attempt to argue that Migdal only related to two-dimensional terrain or that it only related to the download of image data (not elevation data). Rather, Skyline argued that the Examiner's citation to Migdal for the "downloading" element was not correct, and that "there is no disclosure regarding a downloading order based on resolution levels" in Migdal and that this would not have been obvious. *Id.* at GOOG 117-19 (10/5/01 Amendment). In a Final Office Action, the Examiner corrected his citation (noting that data blocks in Migdal were downloaded from a "mass storage device 208"), and again rejected the claims based on Migdal and based on his finding that it would have been obvious to reverse the download order. *Id.* at GOOG 131-32 & GOOG 137 (11/27/01 Final Office Action).

111. Skyline then cancelled many of its claims and amended others. Skyline asked the Examiner to reconsider his obviousness determination, arguing that Migdal taught away from a

download order of lowest resolution to highest resolution, and that Migdal's highest resolution to lowest resolution download order was not a mere reversal. *Id.* at GOOG 151-54 (3/1/02 Amendment After Final Rejection). Again, Skyline did not argue that Migdal was only an image system or that it did not have a remote server. The Examiner thereafter allowed the amended claims.

112. A person of ordinary skill in the art would understand from this prosecution history that the Examiner of the '189 patent application believed that Migdal disclosed all of the limitations of claims 1 and 12, except a downloading order based on resolution levels. The claims were allowed based on Skyline's argument that it would not have been obvious from Migdal alone to reverse the download order.

113. It is my understanding that any prior art considered by the Examiner during the prosecution of a patent application is cited on the face of the patent. The '189 patent cites the Migdal patent, but it does not cite the Cosman article. Thus, it is my understanding that the Examiner for the '189 patent application never considered Cosman, or the combination of Migdal with Cosman.

114. In my opinion, a person of ordinary skill in the art would have been motivated to combine the Migdal patent with the Cosman article. The suggestion or motivation for combining these references can be found in the references themselves: Cosman was cited on the face of the Migdal patent. In addition, Cosman and Migdal both relate to the field of computer graphics and are both directed to solving the problem of rapidly providing texture data for displaying an image.

A. Obviousness

1. *All Asserted Claims: "providing data blocks describing three-dimensional terrain"*

115. The asserted claims recite a method of or apparatus for "providing data blocks describing three-dimensional terrain data." The Court has construed "data block describing three-dimensional terrain" to mean "a block or collection of data or digital information that represents or describes a section of three-dimensional terrain at a particular resolution level and that includes any additional data overlaid on the digital image of the terrain, such as altitude, labels or optional objects." Chang Decl., Ex. 4 at 9-12.

116. The Migdal patent discloses "a method of providing data blocks describing three-dimensional terrain to a renderer." In the Abstract of the Migdal patent, the inventors state that they are disclosing "an apparatus and method for … providing texel [texture element] data relevant for displaying a textured image." Haight Decl., Ex. 52 ('783 patent at Abstract). Furthermore, the Migdal patent provides that "three-dimensional texture data can be used." *Id.* at col. 3:18-19). The Migdal patent also describes a processor (202) and graphics subsystem (220) which "maps texture data … to pixel data in the screen space." *Id.* at col. 6:65-7:21. That is, the texel data, or "data blocks," are provided to a "renderer." The "data blocks" can be "data blocks describing three-dimensional terrain data" because the texture data can represent terrain: "A large amount of texture source data, such as photographic terrain texture, is stored as a two-dimensional or three-dimensional texture MIP-map…. Only a relatively small clip-map representing selected portions of the complete texture MIP-map is loaded into faster, more expensive memory." *Id.* at Abstract.

2. *All Asserted Claims: "renderer"*

117. The asserted claims further indicate that the "data blocks describing three-

dimensional terrain" are provided to a "renderer." The Court has construed the "renderer" to be a "software and/or hardware object that performs at least the following functions: (1) determining and providing to another object the required coordinates in the terrain along with a respective resolution level; (2) receiving the data blocks corresponding to the specified coordinates; and (3) using the received data blocks to display a three-dimensional image." Chang Decl., Ex. 4 at 26-32.

118. Moreover, to the extent that Skyline seeks to construe the term "renderer" more broadly than the Court's construction in an attempt to capture Google Earth, then in my opinion, Migdal also disclosed a "renderer." In particular, in the Migdal patent, processor 202 determines the new tiles that are needed "when the eyepoint and/or field of view changes to ensure that clip-map 540 contains the texel data which is most likely to be rendered for display", processor 202 receives the new tiles from mass storage device 208 and "passes" them to texture memory 226, and graphics subsystem 220 renders an image based on the texture data in texture memory 226. Haight Decl., Ex. 52 ('783 patent at col. 10:42-52, col. 6:55-57, & col. 6:65-7:21, respectively).

3. *All Asserted Claims: "data blocks belonging to a hierarchical structure"*

119. The asserted claims further recite "data blocks belonging to a hierarchical structure which includes blocks at a plurality of different resolution levels." The Court has construed "data blocks belonging to a hierarchical structure" as "data blocks that are organized into multiple levels of resolution, whereby each level contains data blocks at the same resolution, and each successive level contains data blocks of a higher resolution than those in the preceding level." Chang Decl., Ex. 4 at 12-15.

120. The Migdal patent discloses "data blocks belonging to a hierarchical structure which includes blocks at a plurality of different resolution levels." Haight Decl., Ex. 52 ('783

patent at col. 1:47-49 ("The MIP-map consists of a texture pattern pre-filtered at progressively lower or coarser resolutions and stored in varying levels of detail (LOD) maps. *See, e.g.*, the explanation of conventional texture MIP-mapping in Foley et al., *Computer Graphics Principles and Practice*, Second Edition, Addison-Wesley Publishing Company, Reading, Mass. (1990), pages 742–43 and 826–28 (incorporated by reference herein).")). In particular, the Migdal patent "pertains to an apparatus and method for providing texture by using selected portions of a texture MIP-map." Haight Decl., Ex. 52 at col. 3:6-9.

4. *All Asserted Claims: "receiving from the renderer one or more coordinates in the terrain along with indication of a respective resolution level"*

121. The Court has construed "coordinates in the terrain" as "a set of numerical values that identifies a particular location in the terrain," and "terrain" as "the surface features of an area of land, an object, or a material, including color, elevation, and existing objects or structures on the land, object or material." Chang Decl., Ex. 4 at 17-23. The Court has also clarified that "receiving from the renderer" means "something distinct from the renderer receiving from the renderer." *Id.*, Ex. 5 at 8-10.

122. The Migdal patent specifically discloses such use of "coordinates in the terrain along with indication of a respective resolution level." First, the coordinates in the terrain are computed: "Texture coordinates for each pixel quad are ultimately output from scan conversion module 1030. Normalizer and divider 1040 outputs normalized texture coordinates for the pixel quad. Such scan conversion processing is well-known for both two-dimensional and three-dimensional texture mapping and need not be described in further detail." Haight Decl., Ex. 52 ('783 patent at col. 14:28-33). Next, the respective resolution level is computed: "LOD generation block 1050 first calculates an appropriate level of detail for the pixel quad (or pixel) according to standard LOD generation methods based on the individual pixel size and texel

dimension." *Id.* at col. 14:45-49. The "appropriate level of detail" is the same thing as the "respective resolution level." Thus, the system described in the Migdal patent receives from the "renderer" coordinates along with indication of a respective resolution level

5. All Asserted Claims: "providing the renderer with a first data block which includes data corresponding to the one or more coordinates, from a local memory"

123. The Court has construed "first data block" as "the first data block provided to the renderer from the local memory corresponding to the specified coordinates." Chang Decl., Ex. 4 at 15-17. The Court has construed "local memory" as "memory easily accessible to the user's processor, either because it is physically part of the processor or is attached directly thereto, and distinct from the memory of the remote server from which data must be downloaded." *Id*. at 32-34. The Court has also clarified that "providing the renderer" means "something distinct from the renderer providing to the renderer." *Id*., Ex. 5 at 8-10.

124. The Migdal patent discloses "providing the renderer with a first data block which includes data corresponding to the one or more coordinates, from a local memory." The Migdal patent teaches: "According to the present invention, then, a check is made to determine whether texel data for the pixel quad is included within a tile corresponding to the appropriate level of detail. When a texel is included within a tile at the appropriate level of detail, a LOD value corresponding to this tile is output. Otherwise, a LOD value is output identifying a tile at a lower level of detail which includes a substitute texel." Haight Decl., Ex. 52 ('783 patent at col. 14:45-58). Thus, the "renderer" is given access to the tile in a local memory that contains the desired texel corresponding to the coordinates, though possibly at a coarser level of the hierarchy (i.e., "a first data block which includes data corresponding to the one or more coordinates, from a local memory").

6. All Asserted Claims: "downloading from a remote server one or more additional data blocks at a resolution level higher than the resolution level of the first data block which include data corresponding to the one or more coordinates if the provided block from the local memory is not at the indicated resolution level"

125. The Court has construed "downloading" to mean "requesting over a network from a separate computer and receiving on a local computer." Chang Decl., Ex. 5 at 4-8. The Court has construed "downloading … if the provided data block from the local memory is not at the indicated resolution level" to mean "downloading … upon some determination that the block provided from local memory is not at the indicated resolution level." *Id.* at 10-12.

126. The Migdal patent describes how the clip-map is updated by loading the clip-map with relevant portions (data blocks) of the entire MIP-map from a mass storage device. Haight Decl., Ex. 52 ('783 patent at cols. 10:55–11:23). This updating can be conditional on changes in viewpoint: "For example, as the eyepoint shifts, a new row of texel data is added to a tile in the direction of the eyepoint movement. A row located away from a new eyepoint location is discarded. Coarser tiles need not be updated until the eyepoint has moved sufficiently far to require a new row of texel data." *Id.* at col. 10:55–62. Thus, the same change in viewpoint that necessitated the provided block not being at the indicated resolution, causes additional data blocks to be read into the clip-map. The Migdal patent further discloses that a mass storage device containing texture data may be remote: "[u]nder conventional texture mapping techniques, even if texture data were to be accessed from a remote, large texture MIP-map, the rendering of a textured image for display in real-time would be impractical, if not impossible." *Id.* at col. 7:51-58. Thus, the Migdal patent discloses the use of remote locations for storing MIP-map data. Furthermore, the Migdal patent discloses that its system may be connected to a network. *See, e.g., id.* at col. 7:18-21 (stating that the rendered displayed image may be transmitted over a network). Therefore, it would be obvious to someone of ordinary skill in the

art to practice the teachings of the Migdal patent by downloading data blocks from a remote MIP-map accessed over a network when the MIP-map would be too large to store on the local machine. For example, the Migdal patent states, "This hierarchical texture mapping storage scheme allows huge texture MIP-maps to be stored rather inexpensively on the mass storage device 208." *Id.* at col. 7:63-65. A person of ordinary skill would also be aware that a mass storage device could be accessed through a network instead of being local to the machine.

127. Cosman discloses that "[t]errain texture is paged independently for each MIP LOD, and this provides a large reduction in total paging load. Preferential paging of the lower LODs ensures a benign paging overload fallback--terrain would then be textured with a lower resolution version of the right texture." Haight Decl., Ex. 28 (Cosman, p. 62 ¶ 2). Thus, in Cosman, lower level of detail ("LOD") blocks are obtained first, so that if the system gets overloaded and cannot obtain the higher LOD blocks needed for an image in time, the terrain can still be textured with a lower resolution version of the correct texture.

128. The Migdal patent discloses loading a relatively small clip-map representing selected portions of a complete texture MIP-map into a fast memory for rapidly displaying an image. Attempts to access a texture element lying outside of a particular clip-map tile are accommodated by utilizing a substitute texture element obtained from the next coarser resolution clip-map tile that encompasses the desired texture element.

129. One of ordinary skill in the art would have been motivated to combine the technique of Cosman with the system of the Migdal patent in order to ensure that coarser levels of the clip-map are loaded before finer levels of the clip-map when multiple levels of the clip-map must be loaded.

130. Finally, I note that Skyline appears to rely on testimony and documents from

Michael Jones for its assertion that Migdal and Cosman do not render obvious claims 1 and 12. I have reviewed this evidence, and do not agree with Skyline's conclusions. In particular, Mr. Jones states that he did not perform an invalidity analysis or reach a definite opinion regarding whether Migdal or Cosman invalidates the claims of the '189 patent. Moreover, Mr. Jones does note that there is significant overlap between the prior art and the '189 patent.

I declare under penalty of perjury under the laws of the United States that the foregoing is true and correct. This declaration is executed this 2nd day of February, 2007, in Ballarat, Victoria, Australia.

Steven K. Feiner, Ph.D.

Certificate of Service

I hereby certify that, on February 2, 2007, I caused a true and accurate copy of the foregoing document to be served upon all counsel of record for each party by complying with this Court's Administrative Procedures for Electronic Case Filing.

By: /s/ Darryl M. Woo
Darryl M. Woo